



A Proposed Modeling Protocol for Evaluating Information Attacks

by John H. Brand II

ARL-TN-112

January 1999

19990421 042

The findings in this report are not to be construed as an official Department of the Army position unless so designated by other authorized documents.

Citation of manufacturer's or trade names does not constitute an official endorsement or approval of the use thereof.

Destroy this report when it is no longer needed. Do not return it to the originator.

Army Research Laboratory

Aberdeen Proving Ground, MD 21005-5067

ARL-TN-112

January 1999

A Proposed Modeling Protocol for Evaluating Information Attacks

John H. Brand II

Information Science and Technology Directorate, ARL

Abstract

The essence of an information attack is to alter, either by intrusion into and manipulation of a database or by deception, the scenario under which a target mind or organization evaluates and selects future courses of action. The aim is to influence the actions of the target. The method is alteration of the perceived desirability or expected payoff of specific courses of action. This alteration of the information in possession of the target can be described as alteration of the perceived reality under which the target operates. Probable success by an attacker in altering the target's perceived behavior, given a successful manipulation of the target's information, has, in the past, been subjective. A modeling protocol based on the use of game theory is proposed that may, in certain cases, allow optimization of the scenario, or reality, imposed on the target to force the choice of a desired course of action. It should also allow a quantitative estimate of the likelihood of the target's adopting a given course of action. This tool can be used to estimate friendly susceptibility to information attack.

Table of Contents

	<u>Page</u>
List of Figures	v
1. Problem	1
2. Access	2
3. Decisions	4
4. The U.S. Army Decision Process	5
5. Analysis of the Effects of the Attack	10
6. Further Work—The Next Step	20
7. References	23
Distribution List	25
Report Documentation Page	27

INTENTIONALLY LEFT BLANK.

List of Figures

<u>Figure</u>	<u>Page</u>
1. Sample Decision Matrix (From U.S. Army Command and General Staff College [1989], Figure 4-7, p. 4-16)	6
2. Sample Decision Matrix (From U.S. Department of the Army [1997], Figure 5-11, p. 5-25)	7
3. Illustration of Factors Entering a Decision	8
4. Manipulation and Analysis of the Target's Perceived Realities and Sets of Courses of Action for Those Realities	13
5. Determination of the Target Payoff Matrix	14
6. Target and Attacker Payoff Matrices	15
7. Manipulation of the Analysis of Courses of Action	16
8. Determination of the Target's Strategy Mix	17
9. Determination of the Attacker's Best Strategy Mix and Value of the Game	18

INTENTIONALLY LEFT BLANK.

1. Problem

An information attack is a deliberate attempt to alter the reality perceived by an enemy. In the past, the means used to attack a target's information base were deception and persuasion. Persuasion is as old as mankind. Military deception also has an ancient pedigree. The Hittites deceived Ramses II outside the walls of Kadesh around 1286 BC and very nearly destroyed the Egyptian army. The deception means used in that operation are not so very different from the means used in the deception operations now routinely embodied in U.S. operations planning or embodied in the Soviet doctrine of *maskirovka*.*

Deception is the progenitor of a more modern form of information attack, the deliberate and direct alteration of databases, and, hence, perceived reality, by unauthorized entry into and manipulation of computer systems. For the purposes of this paper, an *information attack* is assumed to include an attempted alteration of information in a target's possession by any means, including intrusion or deception.

In the past, a key question has been how likely is successful manipulation of a target's actions, given a successful information manipulation (information attack). This has largely been subjective. The means of coping with the quantitative uncertainty in information manipulation has been the use of whatever resources are available to make the desired course of action look so attractive it is irresistible. A more quantitative means of analyzing and evaluating information attacks may now be available. This is due largely to the increased use of computer-based decision and evaluation aids. These decision aids, such as combat simulations, produce relatively reproducible quantitative results and are driven by scenarios, or representations of reality, that are, in general, governed by information or databases that are numerical in nature. If the input data can be deduced, stolen, or forced, and the decision aids can be obtained, the numerical bases for decisions may be duplicated.

* *Maskirovka* is an elegant Soviet doctrine for controlling the actions of an opponent. It embodies what we call "deception," as well as electronic warfare and more. It is described in Glantz (1989). The historical material is related to the present (then, 1989) in the last chapter. It is worth remembering that the Soviets trained people in their doctrine who we will continue to encounter for a generation to come. For that matter, the Russian military seems unlikely to abandon this elegant and successful doctrine.

Once the numerical bases for decisions are produced by the decision aids—gains or losses from adoption of some given course of action—the decision process may in some cases be emulated. The problem analyzed is the likelihood of successful manipulation of the target's actions, given the alteration of the information in the target's possession. In other words, will the target do as the attacker wishes?

2. Access

Analysis of the results of an information attack assumes the attacker has a great deal of access to a target's information and decision tools. It may be argued that one cannot just assume access to the target's decision process and information base, and in general this is true. However, a successful information attack by intrusion into a computer system, whether by means of a network or by a file manipulation through a delayed action attack such as "chipping" or a previously seeded virus, requires access in order to occur at all.* In general, once in a computer system, the attacker has a great deal of freedom of action, and most intrusions are not detected at the time, if ever. Given that a successful intrusion occurs at all, thus permitting an information attack, the access of the attacker is likely to be very great. Because of this, the access to the information necessary for analysis of the effect of information attack by intrusion is assumed in this analysis as well.

The method of analysis may also apply to many deception operations, as well as intrusions. Although deception operations can be launched with little or no access to a target's information base or reasoning process, the majority of successful ones appears to have been conducted with some, usually much, access. Additionally, the most successful operations have been conducted through reinforcing the preconceived expectation of the target, which required access to the thought processes of the target to set the goal and to the reasoning processes of the target in order to reinforce

* "Chipping," or building malicious code or circuitry into a chip, is not a new idea. Arthur C. Clarke described it in a story, "The Pacifist," in 1956 about a piece of circuitry surreptitiously inserted into a military tactical planning computer. The story has been anthologized in *Tales from the White Hart* (Clarke 1957).

that notion.* It seems reasonable that some deception operations will involve the kind of access to information and knowledge of the target to allow predictive methods to be used.

It can, therefore, be argued that in order for an information attack to be prosecuted successfully access is a *sine qua non*. Without just such access, the attack will probably not occur at all. Therefore, if the attack occurs at all, the conditions required for a predictive protocol may well, but not certainly, be in place.

Assuming access, a protocol for evaluating the effect of information attacks may be formed from the union of tactical combat modeling tools and classical operations research methods for selecting the optimum course of action among several courses of action based on the expected gain or loss for each. Interestingly, the more elaborate the analytical tools for decision making available to a force (e.g., U.S. Army) are, the more predictable the actions of the force may be.

Interestingly, an evaluation protocol may only be useful given an information attack. In the past, use of classical operations research methods to emulate decision processes has often been difficult or impossible. There are simply too many free variables in the usual decision process. A good example is shown in the discussion that follows, which illustrates a decision process using the command estimate process. In the context of an information attack, however, the perceived situation is simplified by the attacker.

No information attack is conducted with the aim of leaving freedom of choice to the target. The target of an information attack is not left with options; the intent is to channel the target's actions. The target is intended to either follow one course of action only, or to be paralyzed, following none. This implies the manipulation of information so that either one course of action is overwhelmingly

* This is demonstrated in the marvelous study, "Deception Maxims: Fact and Folklore" (Central Intelligence Agency 1980). As discussed in this study, an attack through conventional deception is most successful if it aligns with the target's preconceived ideas and expectations. It is also worth remembering that the great deceptions of WW II and Vietnam were conducted with inside information from agents and broken codes in the former, and the news media and hostile intelligence services in the latter. These allowed both information for initial planning and also constant feedback on the success of the operation.

attractive, even if only by comparison, or no course of action is preferable to others, and none look attractive at all. This simplification may allow the use of mathematical methods of evaluating the likelihood of the target's adopting the desired course of action that are impractical in the general case.*

3. Decisions

Evaluating the most probable outcome of a successful information attack is difficult because it involves evaluation of the outcome of manipulation of perceptions on a decision. As stated previously, this study assumes that the information has already been successfully transferred or altered, and the goal is to assess the actual impact of the *successful information manipulation* on the manipulation of the target's *actions*. As stated previously, the information manipulation can be by deception, intrusion, or some other means.

For situations where a logical analysis is used by a target decision-making entity to determine courses of action, the application of game theory or other related disciplines such as linear programming might lead to predictive insight.[†] The prototypical case would be a military force led by individuals with formal training in some analytical staff methodology, such as U.S. battle staffs. Likewise, an attacker's own trained general staff might be increasingly vulnerable to being foxed in turn. The method might be open, with reservations, to application to economic cases as well.

* In the general case, evaluating the probable reactions of a target, given little or no knowledge by the attacker of the target's knowledge of the situation, and in ignorance of the target's values and thought processes, is a task best left to intuition. Unfortunately, in many cases this is the only real avenue. For many people, especially the violent personalities likely to seek, gain, and maintain power in the Third World, emotion is king. Reason is at best a distantly related handmaiden. Curiously, cases that are subject to analysis seem more likely to occur in targets that are better educated and better trained than in cases such as a criminal without formal training, or an emergent popular leader. Many "popular" leaders, of course, are quite rigorously trained by some hostile nation state.

[†] Zero-sum game theory and linear programming (LP) problems may be expressed in either formalism. This is extremely convenient, as LP software is easily available. The relationship between the two formalisms is described in Bennion (1960).

4. The U.S. Army Decision Process

The formalism used by U.S. Army battle staffs is illustrated in two examples of a sample decision table from the G3 (operations) planning process. The first (Figure 1) is taken from the Command and General Staff College (CGSC) text on the command estimate process (U.S. Army Command and General Staff College 1989).^{*} The second (Figure 2) is reproduced from the newly revised manual on staff operations (U.S. Department of the Army 1997).[†] The “school solution” illustrated in the text is paper-and-pencil based; field units now use sophisticated combat and logistical simulations, but the basic logic applies. Several observations are in order.

In this process, a mission statement leads ultimately to the determination of possible courses of action: “attack with *three* battalions here, defend with one *here*, a barrier *here* ...” These are formulated, evaluated by war gaming, and the decision is made in the light of war game results conducted under a given scenario, or perceived reality. In the general case, individually chosen and weighted factors enter the decision process. This is illustrated in Figure 3. If that reality is tailored, then the determination of possible courses of action and the estimate of the value of the courses of action will be altered also. The perceived relative merit of a course of action is thus determined, based on the factors considered important by the commander. As can be seen in the example, which is only one of many staff decision matrices, without simplification or channeling of the decision, the judgmental factors make predictive methodologies themselves highly judgmental, unless the basic perception of the situation is manipulated or accessed in some way. That is, however, the essence of an information attack.

In the absence of manipulation, the analysis and selection of courses of action, the war gaming of the courses of action, and the formulation of the decision matrix are highly subjective and full of free variables. The process must be simplified or channeled in some way. This can be by

^{*} With the embodiment of the methodology of this ST in the latest revision of FM 101-5 (U.S. Department of the Army 1997), the text will presumably be revised.

[†] The May 1984 edition did not incorporate this excellent material.

CRITERIA ¹	WTS ²	COURSES OF ACTION ³		
		1	2	3
Simplicity	2	2 ¹	1	4
Surprise	3	1	3	2
Speed	5	2	10	5
Mass	1	1	2	1
Combined Arms	1	3	4	3
Security	1	1	2	2
CSS	2	2	3	4
Objective	1	4	4	4
C ⁴	1	3	2	3
Offensive	1	2	2	1
Total		21 ⁴	26	29
Weighted total		35 ⁵	48	61

Note that the subsidiary factors do not *dominate* the decision: all the courses of action have *decent* C2, etc. The driving factors in this case are two: "surprise" and "speed." Furthermore, if, for instance, an attacker left "surprise" out of a prediction of the target's actions, the *relative* rank based on "speed" alone would not change from the example shown. Decisions based on "speed" are driven by factors such as Tables 4-2 and 4-3, "Unopposed rates of movement" --predictable.

¹Criteria are any factors that pertain to the mission. Options include: battlefield operating systems, trends of Airland Battle, ODOMA, critical events. They may be assigned by either the commander or staff. If the criteria are qualitatively the same for each course of action, they may not need to be displayed.

²Courses of action are those that are selected for war gaming.

³The principal staff officers assign numerical values for each criterion after the courses of action are war gamed. These values reflect the relative advantages or disadvantages of each criterion for each course of action. In the example above, course 3 is clearly the best.

⁴The numbers are totaled to provide a subjective evaluation of the best course of action without weighting one criterion over another.

⁵Should the commander desire to emphasize one criterion as more important than another, he assigns weights to each criterion based on relative importance.

The weights are multiplied by the initially assigned score in each column.

The scores are totaled to provide a "best" course of action based on weights assigned by the commander.

Figure 1. Sample Decision Matrix (From U.S. Army Command and General Staff College [1989], Figure 4-7, p. 4-16).

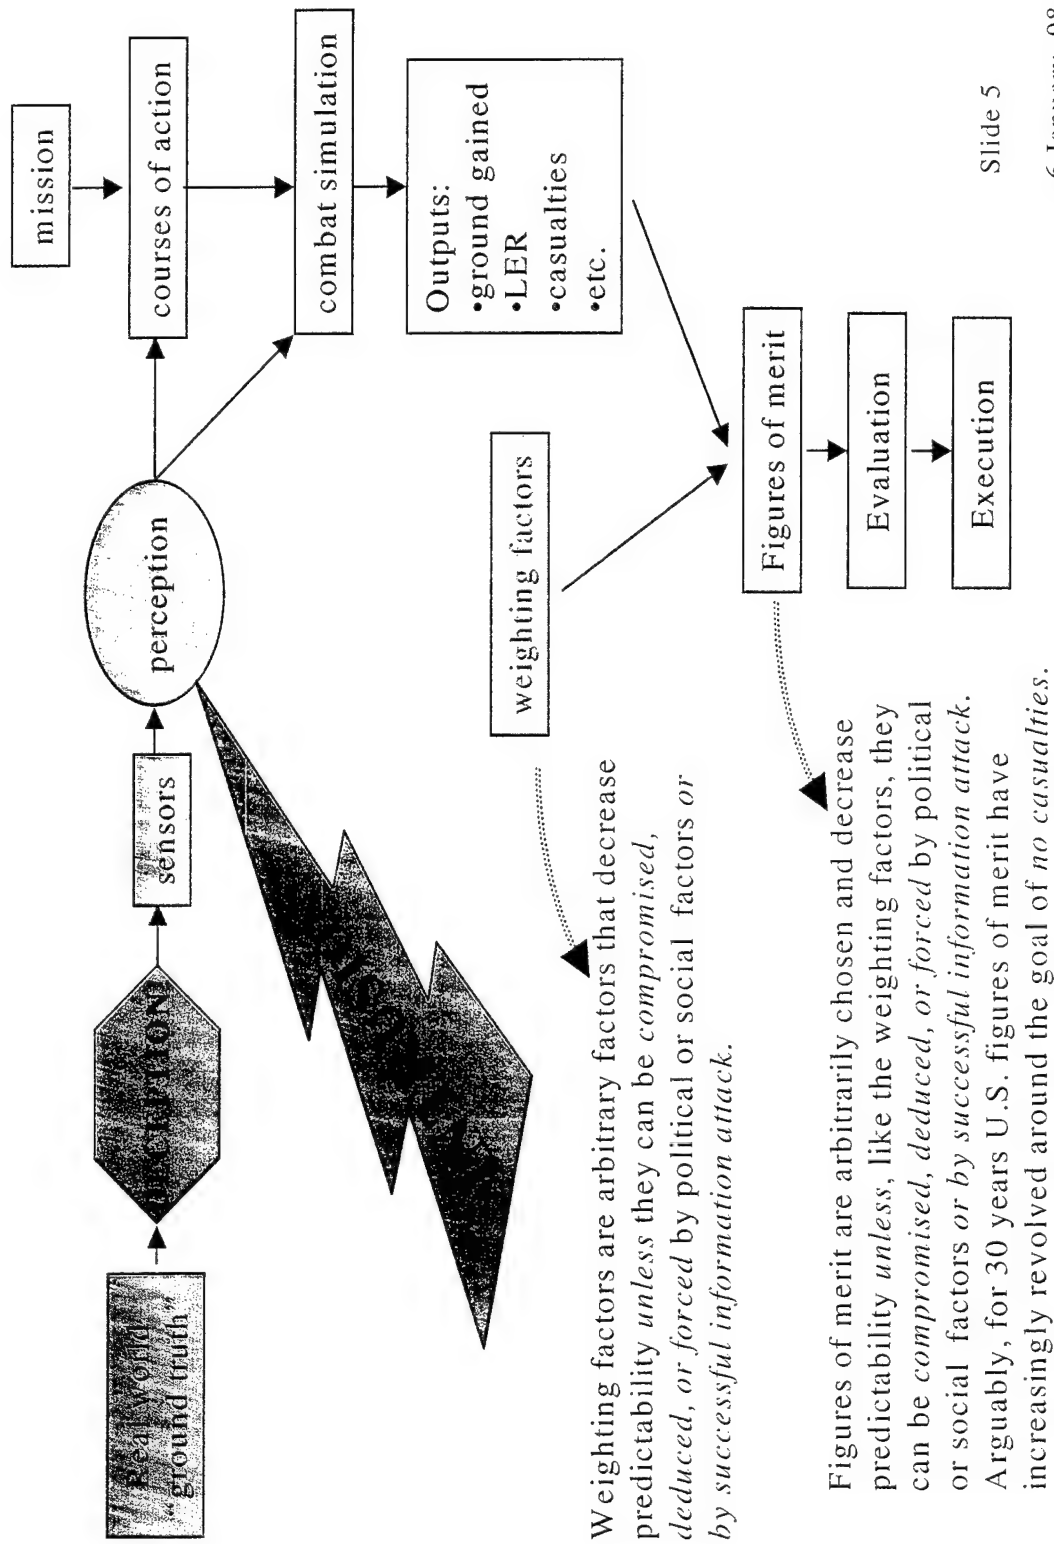
CRITERIA (note 1)	WT (note 2)	COA 1 (note 3)	COA 2 (note 3)	COA 3 (note 3)
Maneuver	3	2 (6)	3 (9)	1 (3)
Simplicity	3	3 (9)	1 (3)	2 (6)
Fires	4	2 (8)	1 (4)	3 (12)
Intelligence	1	3 (3)	2 (2)	1 (1)
ADA	1	1 (1)	3 (3)	2 (2)
Mobility/ Survivability	1	3 (3)	2 (2)	1 (1)
CSS	1	2 (2)	1 (1)	3 (3)
C ²	1	1 (1)	2 (2)	3 (3)
Residual Risk	2	1 (2)	2 (4)	3 (6)
C ² W	1	2 (2)	1 (1)	3 (3)
TOTAL Weighted TOTAL		20 (37)	18 (31)	22 (40)

NOTES:

1. Criteria are those assigned in Step 5 of the war-gaming process.
2. Should the CofS/XO desire to emphasize one as more important than another, he assigns weights to each criterion based on relative importance.
3. Courses of action are those selected for war gaming.

Procedure: The staff assigns numerical values for each criterion after war-gaming the COA. Values reflect the relative advantages or disadvantages of each criterion for each COA action. The lowest number is best. The initially assigned score in each column is multiplied by the weight and the product put in parenthesis in the column. When using weighted value, the lower value assigned indicates the best option. The numbers are totaled to provide a subjective evaluation of the best COA without weighing one criterion over another. The scores are then totaled to provide "best" (lowest number value) COA based on weights the commander assigns. Although the lowest value denotes the best solution, the best solution may be more subjective than the objective numbers indicate. The matrix must be examined for sensitivity. Although COA 2 is the "best" COA, it may not be supportable from a CSS standpoint. The decision maker must either determine if he can acquire additional support or if he must alter or delete the COA.

Figure 2. Sample Decision Matrix (From U.S. Department of the Army [1997], Figure 5-11, p. 5-25).



Slide 5

6 January 98

Figure 3. Illustration of Factors Entering a Decision.

external pressure (e.g., the media) or by compromise of the decision algorithm. The judgment factors are then transparent to the attacker. The selection of the courses of action in the general case, therefore, is not amenable to analysis with game theory or its cousins, though a judgment evaluation may be aided considerably by combat simulations run and evaluated with the target's decision methodology.

The decision matrix from FM 101-5 (U.S. Department of the Army 1997) is now a rank-ordered rating.* This may decrease the impact of initial conditions in the combat simulation used to generate figures of merit. Note that in this example, as in the previous case, a factor dominates the decision. This cannot be relied on in the general, unmanipulated case, even though present in the example. An additional observation about this example is that, if the miscellaneous nondominant factors such as combat service support are ignored in this matrix, the decision does not change. This is because, in general, bad plans are not made deliberately. Good alternative courses of action are all supportable, all have decent fire support, etc.

If real battle staffs perform like the textbook examples, some key attributes will likely be the deciding factor, picked out by the weighting functions. In fact, it can be argued that U.S. planning has become more dependent over the last 30 years on a single, overriding imperative: low casualties. If two or more factors should dominate, however, the protocol can also be followed for each factor, and the course of action selection by the target estimated based on the agreement of the protocol for all the factors for a single course of action. If the answer does not converge in this way, the assumed manipulated reality is inadequate to channel actions with any certainty. The scenario should be altered and the analysis performed again to find an unambiguous answer.

* In passing, the new school solution with forced choice rank ordered rating may meet some resistance. The method artificially enhances differences between choices of action where there may, in fact, be little or no real difference. It also de-emphasizes differences where there may be substantial differences. For instance, two courses of action may be identically supportable in terms of ton-miles, but one is forced to be rated better than another (1 vs. 2). One may produce three times the casualties of another, but the better is rated as merely better (1 vs. 2) rather than three times as good as the other. The older method accounted for these problems.

The most important thing to glean from the discussion of gaming of courses of action in the text is that the gaming approach outlined in ST 100-9 (U. S. Army Command and General Staff College 1989) is essentially a zero-sum decision. That is, one side gains according to some measure; the other side loses. For example, the measure may be ground gained or casualties. One side gains ground; the other loses it. One side loses men and vehicles; another side gains kills, and so on.

Mathematically, the application of non-zero-sum theory is more appropriate, but as of now, the doctrinal approach is simpler. Perhaps with the increasing use of more complicated simulation tools, that will change. The attacker is, of course, not only well aware of the difference in goals and the disparity of values, but in this scenario has acquired the necessary predictive tools to generate them. Non-zero-sum game theory is thus essential for the attacker's decision, but not needed for the attacker's evaluation of the target's decisions.

It should be noted that, in the absence of an information attack or detailed inside information, the problem may still be analyzed, but with less confidence. With only the information in either of these two example decision matrices, a prediction can be made, but game theory is not adequate for the prediction. The necessary payoff matrix cannot be filled out. The decision process can likely be emulated, however, if the combat simulation used by the target to calculate measures of effectiveness can be used to fill out a likely payoff matrix that the attacker feels is consistent with the target's values and past judgement. The intermediate steps in the war-gaming process are necessary. Given those, a payoff matrix the attacker feels is similar to what the target would generate can be filled out and an estimate made using the formalism of game theory, which should emulate the judgment call of the staff. The fidelity of the emulation cannot be guaranteed, but the process should be at least indicative.

5. Analysis of the Effects of the Attack

As mentioned, the basis of the attack is the manipulation of a perceived reality in such a way that the problem is simplified and analyzed with the appropriate decision aids to generate payoff

matrices. One question that must be addressed is, "How robust are combat simulations under different initial scenarios, or perceived realities?"

Higher-level headquarters have increasingly sophisticated combat and logistic simulation tools available, with output that may or may not be believable in absolute terms. The output is usually regarded as at least consistent. That consistency is important. Although one may not have a lot of confidence in the probable absolute accuracy of a given prediction of a measure such as casualties, the accuracy of the *relative* results of two courses of action seems reasonably well founded by experience with combat simulations. That is, the casualty ratios in two "battles" that are subsequently "fought" in a combat simulation will likely be reasonably accurate, but the actual values of the casualties predicted by a given simulation using historical data as the scenario may differ substantially from the real battle. Thus, at least for the game theory formalism, the *strategy mix* should be robust, but the *value of the game* (wins or losses) will be altered by any multiplicative factor that relates the war game results to real battles.

The strategy mix is the determining factor for prediction of the selection of a course of action, and should be robust within a given simulation, although the value of the game might be less so. That is, an information attack should increase the probability of a course of action so that it dominates the strategy mix, and the others have low probabilities. The value of the game is important in estimating whether the target may be induced to act at all. This is important in analyzing an attack aimed at paralyzing the target. Such an attack should aim at no course of action predominating, and the value of the game being so poor that action is not taken.

The outcome of a given course of action also depends on what an enemy might do. That is, course of action i might have outcomes j depending on how the enemy reacts. This immediately suggests a rectangular payoff matrix, and the formalism of a decision optimization discipline such as game theory or, for the zero-sum case, linear programming. The discussion in this report is based on the formalism of game theory. This is because, although there is an extensive and mature analytical discipline of optimization of courses of action (linear programming is one methodology), the writer is more familiar with the game theoretical formalism than the others, and the game theoretical formalism extends to the non-zero-sum case.

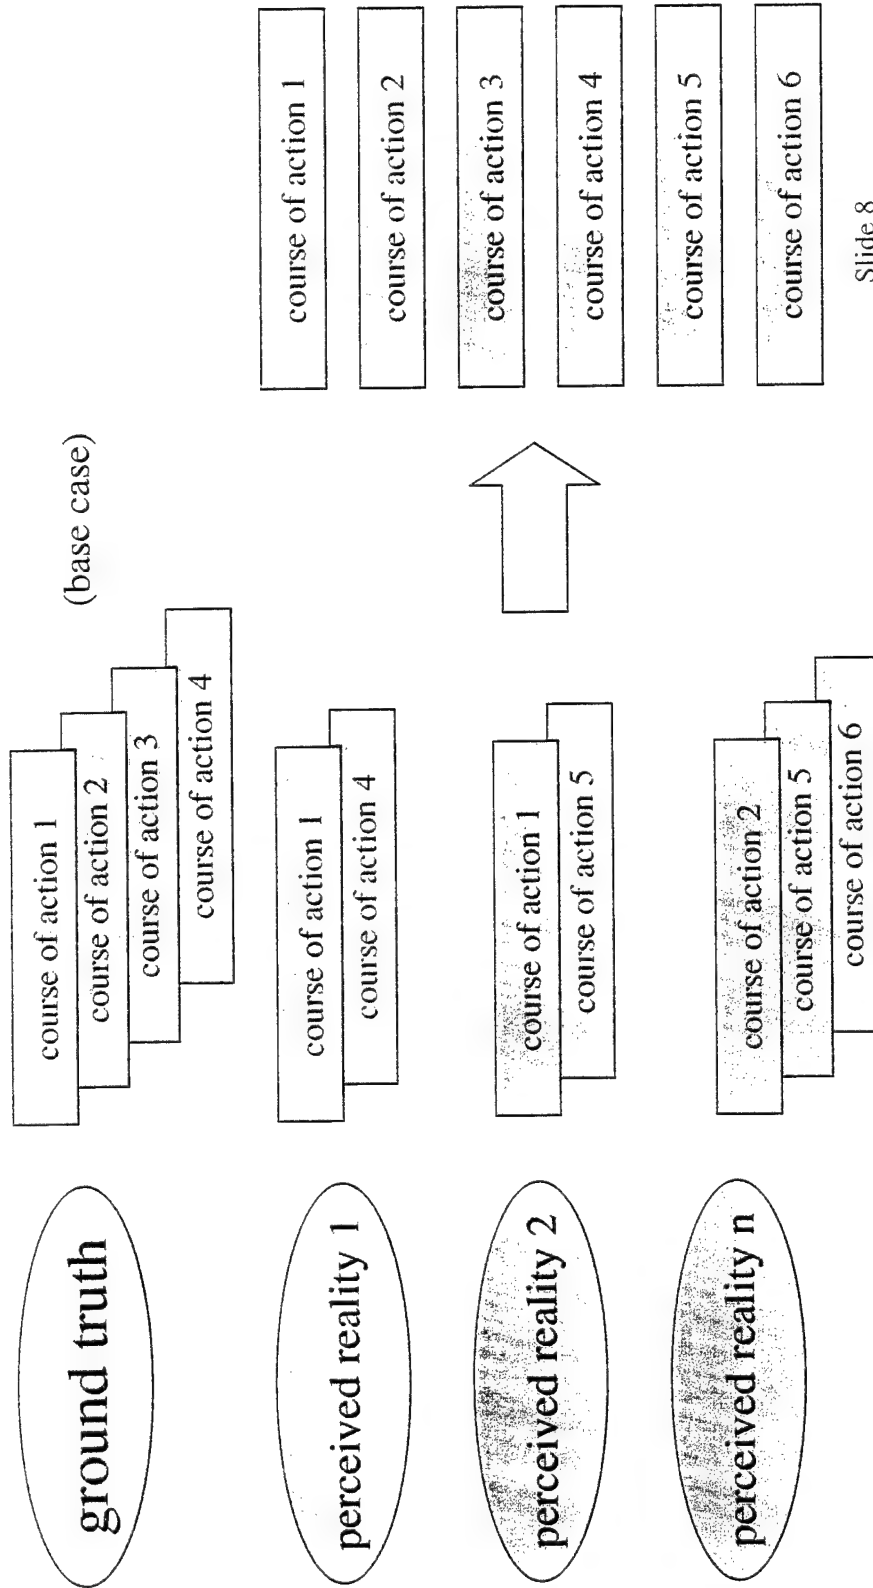
Once a payoff matrix is generated, game theory allows determination of the optimum mix of strategies. That is, for a player playing the same game many times, independently of previous results, choice of strategy can be made randomly within the probabilities in the strategy mix. That is, the formalism of evaluation of an $m \times p$ payoff matrix, corresponding to m friendly courses of action and p enemy courses of action, allows the determination of a series of probabilities for the optimum mix of the m friendly courses of action. If the game were played according to the optimum strategy mix, and the individual strategies in the strategy mix were randomly selected according to those probabilities, the outcome would be optimal. This is also the expected outcome of an evaluation process by a good staff and commander. Indeed, the long-term payoff, the value of the game, can also be determined from the payoff matrix.* As stated previously, the goal of an information attack would be to alter the reality that produces a perceived payoff matrix so that either action is paralyzed (no strategy predominated in the strategy mix, or all look about equally likely, and the value of the game is bad) or a single strategy predominated the mix (action is channeled).

One problem is that an attacker, unless a graduate of the American staff system, is unlikely to evaluate his or her courses of action in the same way or to choose the same figures of merit. That is, his or her payoff matrix may be considerably different. This is especially so if the chosen figure of merit is radically different from U.S. usage; an enemy may be completely indifferent to casualties, or even desire them for their effect on the media or the U.S. intelligentsia. This leads to the necessity for two sets of payoff matrices—one based on the U.S. methodology and figure of merit, and another based on both that payoff matrix and the hostile methodology and figure of merit. This dual-entry matrix, by the way, automatically requires the methods of non-zero-sum game theory.

The generation of the set of target payoff matrices based on the alternate realities that may be forced on the target by the attacker is illustrated in Figures 4–9. This first step is to generate the set of alternative courses of action that the target and attacker might use, under all the proposed perceived realities, shown in Figure 4. The use of the target's evaluation method to determine the attractiveness of a course of action under the supposed conditions of a perceived reality is illustrated

* As discussed previously, a complication is the fact that a given battle or conflict is not truly a zero-sum game, although decisions in combat often are. This simple outline is based on the assumption of a mix of zero-sum and non-zero-sum decisions, depending on by whom and why the decision is made.

Consider a battle staff formulating alternative courses of action and then gaming them to obtain relative value.



Slide 8

20 October 97

Figure 4. Manipulation and Analysis of the Target's Perceived Realities and Sets of Courses of Action for Those Realities.

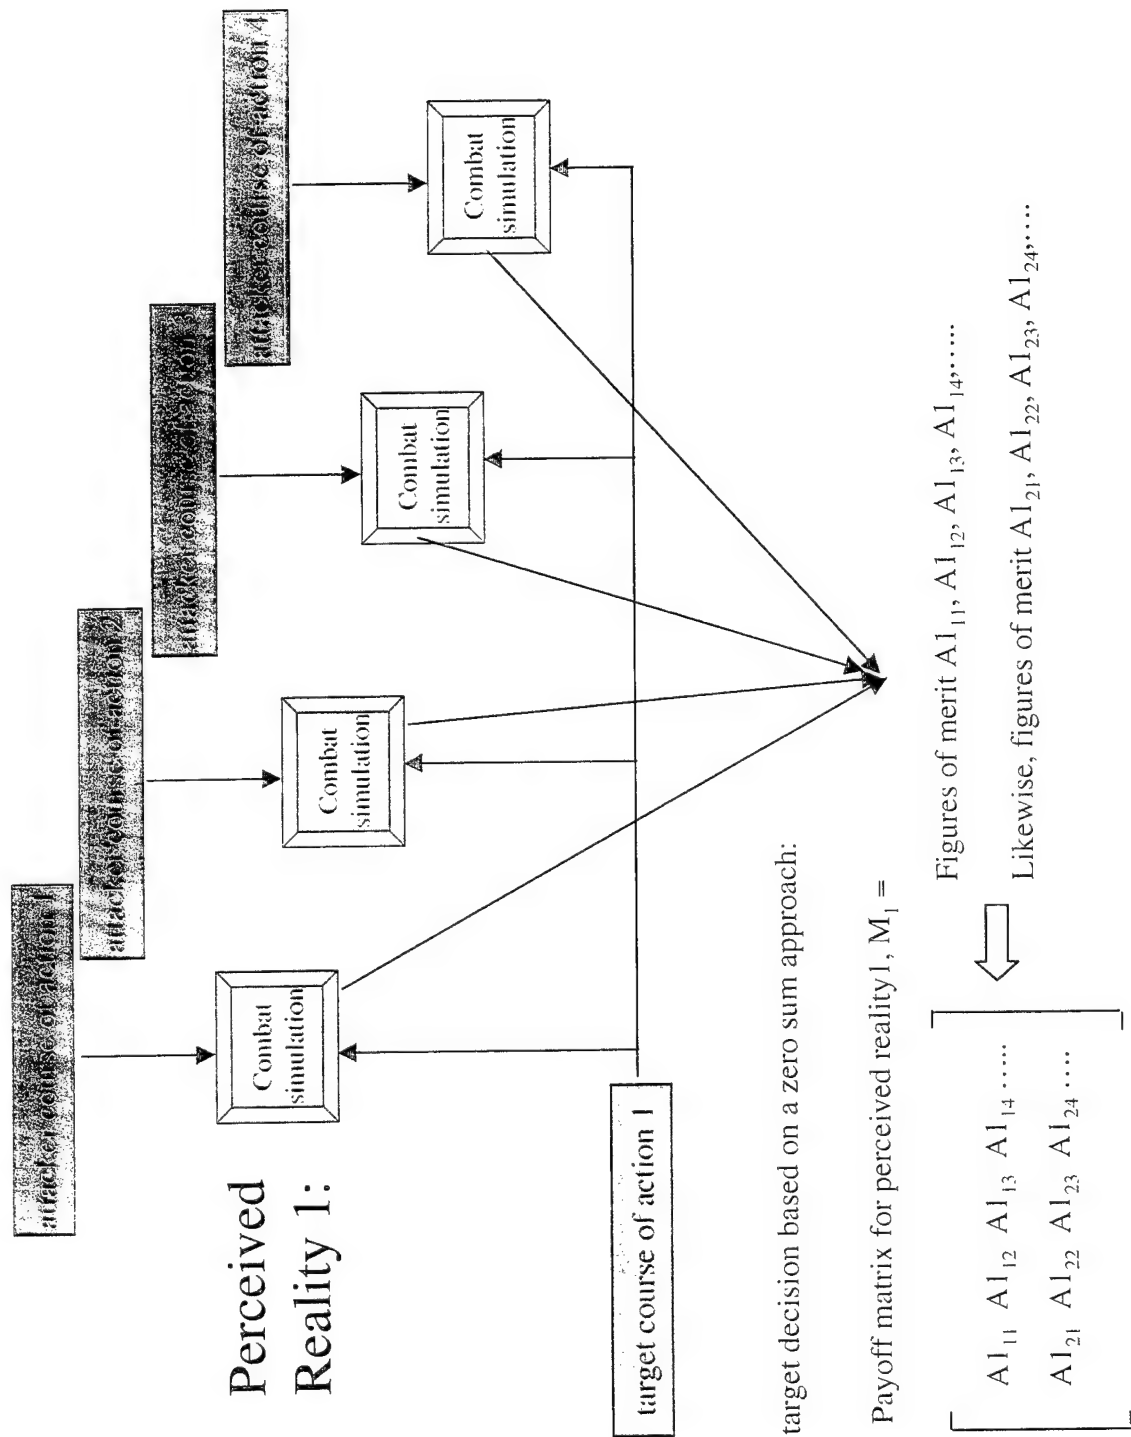


Figure 5. Determination of the Target Payoff Matrix.

Zero sum formulation:

Target's strategy mix:

The target's decision is based on a zero sum approach, in practice

$$\text{Payoff matrix for reality 1, } M_1 = \begin{bmatrix} A1_{11} & A1_{12} & A1_{13} & A1_{14} & \dots \\ A1_{21} & A1_{22} & A1_{23} & A1_{24} & \dots \end{bmatrix} \Rightarrow ST_1 = \begin{bmatrix} P_{11} \\ P_{12} \\ \vdots \\ \vdots \\ \vdots \end{bmatrix}$$

Likewise, for realities 2, 3,...

Non-zero sum formulation:

The attacker must take into consideration both what the attacker feels it gains and also what the target feels it gains. Each will probably use different means of computing the desired factors, and the factors may well be different. Hence, A is the factor valued by the target and B the factor valued by the attacker.

$$\text{Payoff matrix for reality 1, } MM_1 = \begin{bmatrix} A1_{11}, B1_{11} & A1_{12}, B1_{12} & A1_{13}, B1_{13} & A1_{14}, B1_{14} & \dots \\ A2_{11}, B2_{11} & A2_{12}, B2_{12} & A2_{13}, B2_{13} & A2_{14}, B2_{14} & \dots \end{bmatrix}$$

Figure 6. Target and Attacker Payoff Matrices.

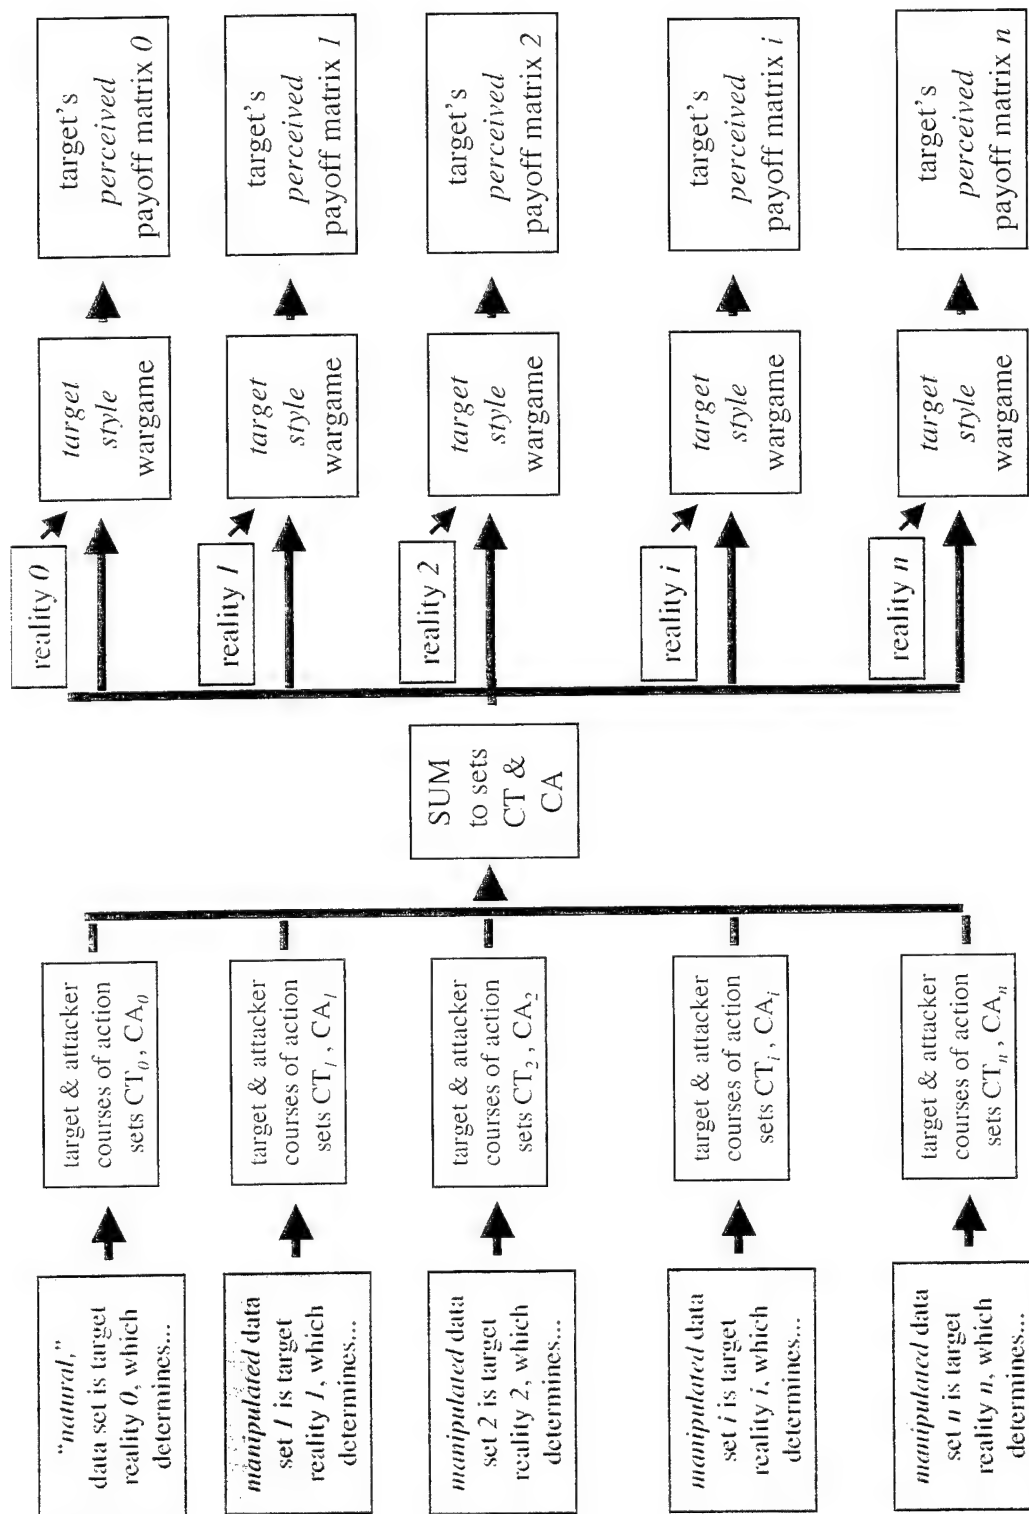


Figure 7. Manipulation of the Analysis of Courses of Action.

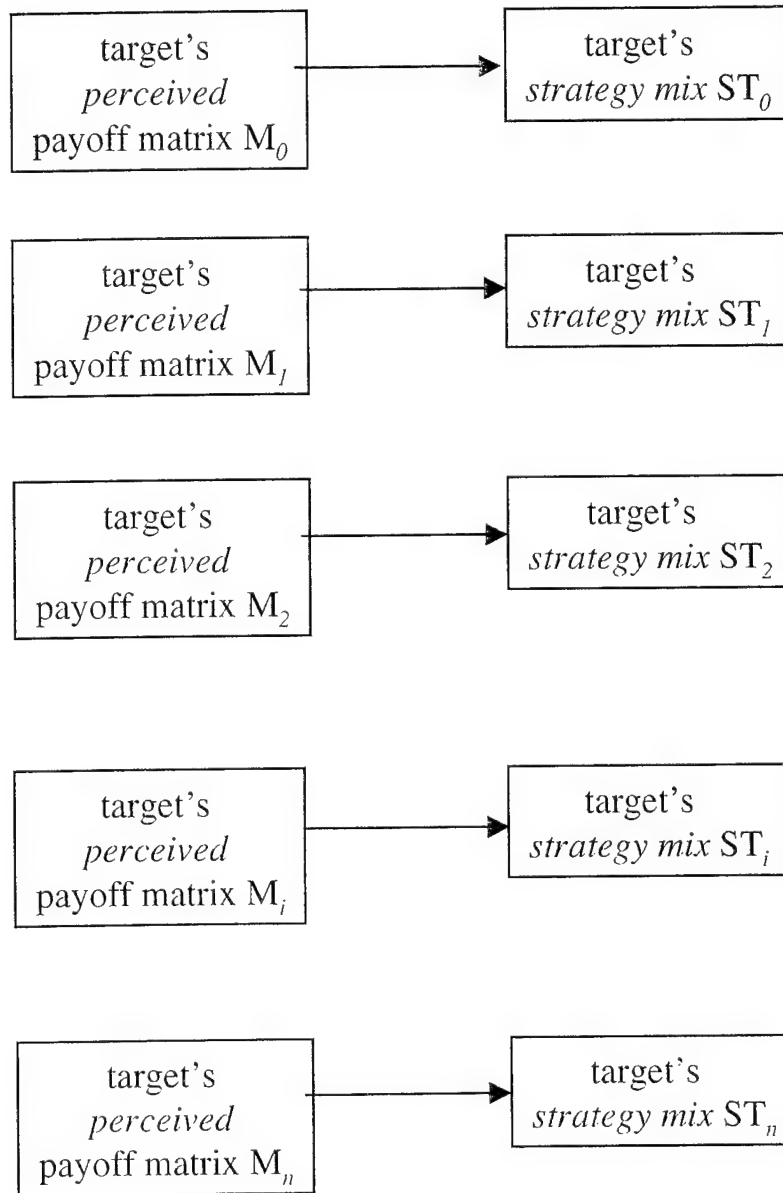


Figure 8. Determination of the Target's Strategy Mix.

in Figure 5. As mentioned previously, if the situation cannot be simplified to a single, dominating factor, the analysis can be run for each factor and the target's selection gauged on the basis of whether the course of action desired by the attacker looks attractive to the target according to each factor. If it does not, the target has an ambiguous choice and the attacker must change the scenario or reality.

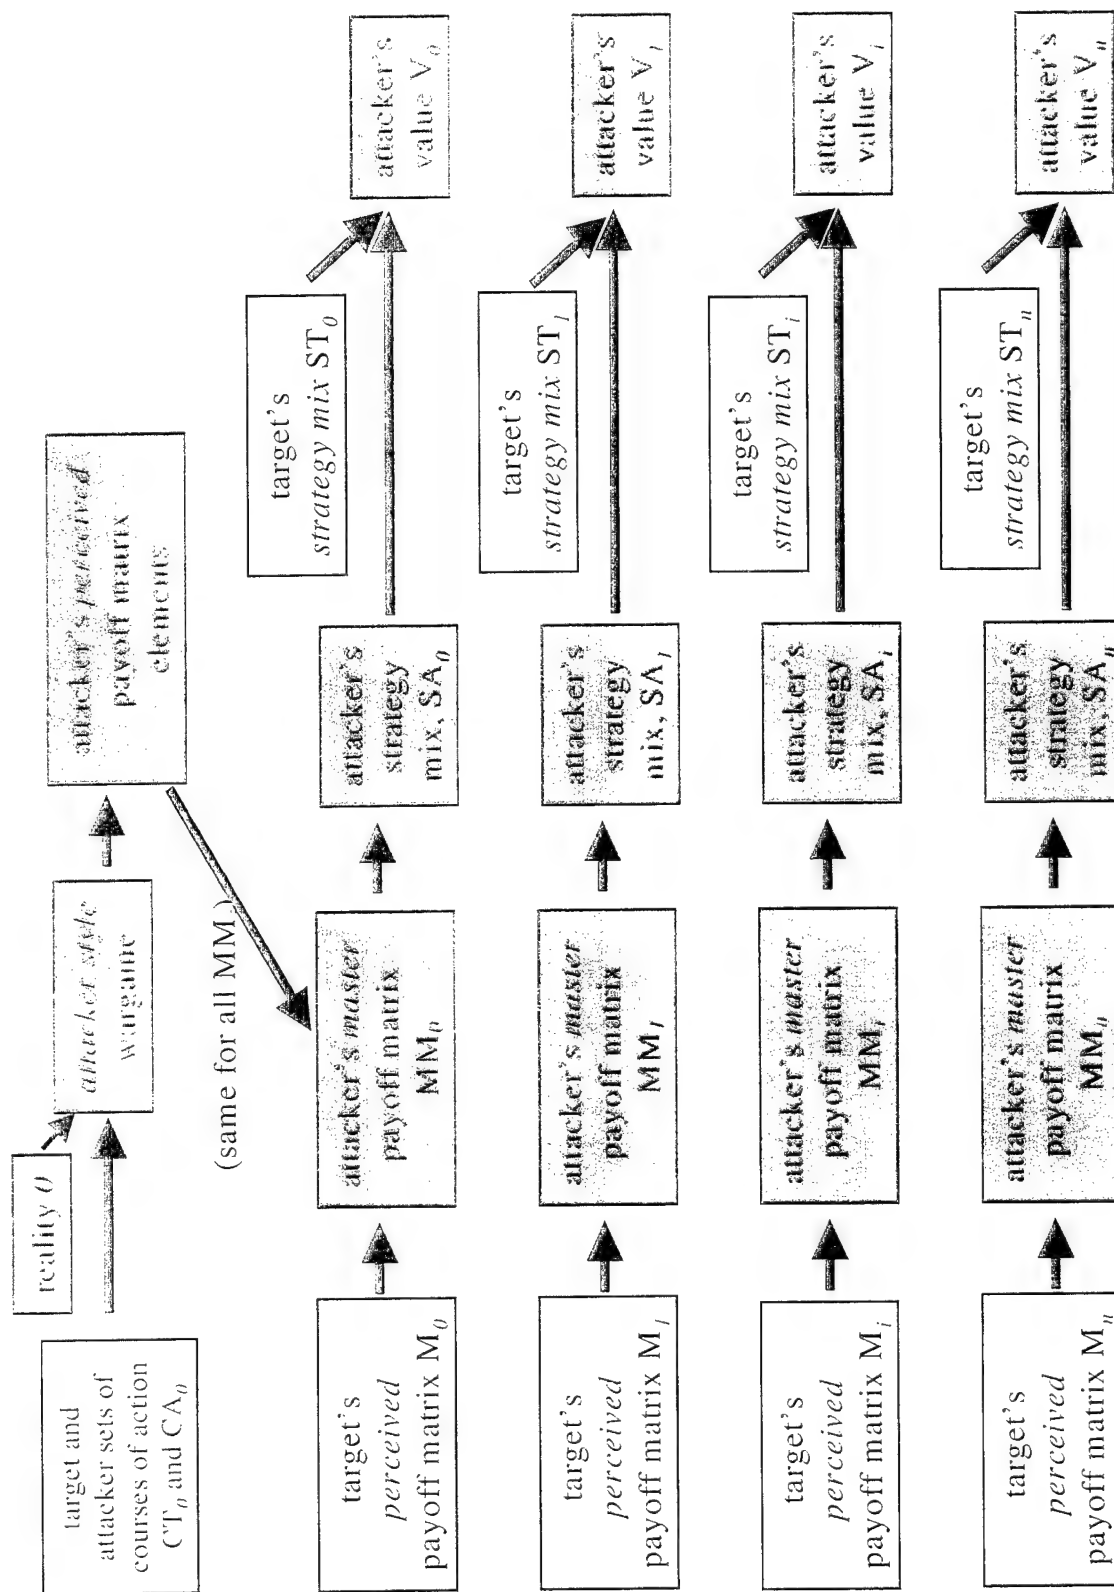


Figure 9. Determination of the Attacker's Best Strategy Mix and Value of the Game.

The different courses of action are war gamed, using the target's war-gaming method, under the conditions pertaining to the perceived alternate views of reality or scenarios chosen by the attacker, to determine what value the target would assign to a given course of action. The result is a series of payoff matrices, one for each scenario or perceived reality.

It seems at this point that the matrix for each perceived reality must include values assigned to each course of action in the consolidated list of courses of action across all perceived realities. If a given course of action really is not appropriate for a given perception of reality, no harm will be done as its payoff would be very poor, and its chance of being invoked in a strategy mix minuscule.

The attacker must then determine the payoff from the attacker's own point of view. This is illustrated in Figure 6. The attacker strategy mix for each possible perceived reality can then be determined from the non-zero-sum payoff matrix formed from the perceived values for target and attacker under that reality. It is important to remember that the target values are based on a target style war game under the conditions of the perceived reality, and the attacker values are generated by use of an attacker style war game based on the ground truth reality. This process is illustrated in Figure 7.

The likelihood of a target staff and commander adopting the desired strategy can then be estimated, the alternate reality adjusted if necessary, and the process run again. This determines the desired form of the alternate reality, facilitating planning of an integrated information attack with, for instance, military deception, manipulation of the press, and possibly covert manipulation of databases. The value to the hostile strategist can then be estimated based on the hostile methodology for evaluating courses of action, with the difference in the uncontaminated reality vs. the adjusted reality clearly delineated. This is illustrated in Figure 8.

Also illustrated in Figure 8, a successful analysis can indicate which reality is most likely to yield the target's adoption of the course of action desired by the attacker. The strategy mix determined for the target by the target's analytical methodology indicates the probability of the target actually adopting a given course of action within that reality. If the desired attacker goal is paralysis, the

strategy mix should end up with all courses of action having about the same probability, and the value of the game, to the target, appearing really poor. If the desired goal to the attacker is for one strategy or course of action to be the most likely choice, the strategy mix will indicate a high probability for that strategy or course of action, and the value of the game will be such as to induce action.

These are likely to be relatively robust as long as the relative outcome between two or more strategies in a combat simulation or other analytical method is reasonably believable, but less confidence can be invested in the computed value of the game. The attacker may then use non-zero-sum game methods to compute the optimum strategy mix, illustrated in Figure 9.

An interesting question poses itself at this point. A reasonable information attack would almost certainly not be a one-shot deal, aimed and fired like an unguided missile. The attacker would very likely monitor the progress of the attack and continue to feed the preconceptions by a deception operation or further intrusion, or both. Historically, at least in deception operations, the target has been fed substantiating information right along. The target will also be systematically looking for confirmatory information as the operation continues by targeting key areas for reconnaissance, and so forth, which should be provided. Whether it is desirable to risk the information attack by further deception or by continuing access through further intrusions and manipulations is an interesting issue that is beyond the scope of this report.

6. Further Work—The Next Step

The protocol is tried out next. For simplicity, a single means of war-gaming is used. This may be the manual war game method or use of the Modular, Semi-Automated Forces (ModSAF) model. ModSAF is highly desirable, but manipulation of the scenario using ModSAF may not be practical. The great merit of ModSAF, however, is that the operator can fight a poor battle under the conditions of ground truth. The poor battle under the ground truth is the battle corresponding to a course of action that seems desirable under a perceived reality. The difficulty lies in changing the scenario

enough to obtain different, distinguishable realities. These variations on the basic scenario are essential for evaluation of the course of action under the presumed conditions of the perceived, but false, realities imposed on the target during a successful information attack. A useful method may be to alter the input tables for a key parameter such as trafficability. This also has the merit of being the kind of alteration an information attack might exploit.

When the zero-sum payoff matrices are generated, the game will be recast as a linear programming problem and the matrix solved using any standard linear programming package. The means of solving the non-zero-sum game is not yet determined. There does not appear to be a general solution (Vorob'ev 1971), although there are approaches that may suffice.* However, the information and method asymmetry in an information attack may allow for a simplification that permits a solution.

The normal discussion of non-zero-sum games assumes equal knowledge of the payoff to both sides, and a common methodology. That is, both are playing the game. In fact, the game is very different for attacker and target. A successful information attack that channels action will reduce the freedom of action of the target to one course of action. This will be a $1 \times p$ game, which should be easier to analyze. This is so because the choice of courses of action by the target is assumed to be based on zero-sum reasoning; hence, the target choice of action can be easily analyzed. The assumed reality is also assumed to be structured so that the choice of courses of action by the target can be assumed to be degraded to a single choice. This yields an attacker payoff matrix with one row. If it is not, the protocol must be repeated with a different set of scenarios, or perceived realities, until the attacker's choice of target strategies dominates in the target strategy mix. The attacker is then free to select his or her own strategy based on maximizing gain for the attacker, or inflicting maximum hurt on the target, or some mix. Note that this essentially converts the two independent payoff values into some utility function, resulting in a simple choice of maximizing or minimizing components of a vector.

* The first try will use the method outlined on pp. 213 and 214 of *The Compleat Strategyst* (Williams 1954), one of a RAND series. One assumes both sides are playing against a neutral mother nature.

If the attacker is indifferent to whether the target chooses one or many courses of action, the problem becomes much more difficult. The strategies can be examined for strict dominance and an equilibrium point sought, but the solution is not clear in the general case. A geometrical approach may work in some cases.*

A further step to aid this analysis might be the construction of a stochastic model of computer network models. This would describe a net in terms of numbers of elements, type of security, and known historical probabilities of successful intrusion through poor security practices, etc. Whether this type of model would be useful is unknown at this time, but it would allow estimation of likelihood of successful intrusion without the use of detailed engineering models, which is the only method now available. A generalization such as this would also allow for historical data on factors such as inadequate training and carelessness, which are difficult to build into an engineering model but are always present.†

It is proposed that a survey of the various means of evaluating courses of action used by friendly and possible enemy decision entities (staffs) be conducted, and, if enough information can be gathered, the result of an information attack be played out for some limited scenario. A sensitivity analysis should be conducted to determine whether this methodology is indeed robust to uncertainties expected in realistic cases. This will require a team approach, the content to be determined by the elaboration of the demonstration involved.

* This procedure is discussed exhaustively in *Games and Decisions* (Luce and Raiffa 1957). The geometrical representation approach mentioned is illustrated on pp. 93ff.

† For example, the choice of passwords can be an obvious weakness. Historically, a large fraction of passwords are chosen in such a way that they are cracked relatively easily. Thus, some proportion of attempted access might be assessed as having encountered a password that can be cracked in a shorter time than in other cases. This likelihood can be also parameterized in terms of the type of net that is attacked: commercial nets might not have password strength checking routines and password aging; tactical nets almost certainly will.

7. References

- U.S. Army Command and General Staff College. *The Command Estimate*. ST 100-9, Fort Leavenworth, KS, July 1989.
- U.S. Department of the Army. *Staff Organization and Operations*. FM 101-5, Washington, DC, May 1997.
- Bennion, E. G. *Elementary Mathematics of Linear Programming and Game Theory*. Ann Arbor, MI: Michigan State University, 1960.
- Central Intelligence Agency. "Deception Maxims: Fact and Folklore." Deception Research Program, Office of Research and Development, Washington, DC, April 1980 (ADB199481).
- Clarke, Arthur C. "The Pacifist." *Tales From the White Hart*, New York: Ballentine Books, 1957.
- Glantz, D. M. *Soviet Military Deception in the Second World War*. London: Frank Cass, 1989.
- Luce, R. D., and H. Raiffa. *Games and Decisions*. New York: John Wiley and Sons, 1957.
- Vorob'ev, N. N. *The Development of Game Theory*. Translated by Erika Schwoediaver, Working Paper No. 2, Department of Economics, New York University, August 1971.
- Williams, J. D. *The Compleat Strategyst*. New York: McGraw Hill, 1954.

INTENTIONALLY LEFT BLANK.

NO. OF
COPIES ORGANIZATION

2 DEFENSE TECHNICAL
INFORMATION CENTER
DTIC DDA
8725 JOHN J KINGMAN RD
STE 0944
FT BELVOIR VA 22060-6218

1 HQDA
DAMO FDQ
D SCHMIDT
400 ARMY PENTAGON
WASHINGTON DC 20310-0460

1 OSD
OUSD(A&T)/ODDDR&E(R)
R J TREW
THE PENTAGON
WASHINGTON DC 20301-7100

1 DPTY CG FOR RDE HQ
US ARMY MATERIEL CMD
AMCRD
MG CALDWELL
5001 EISENHOWER AVE
ALEXANDRIA VA 22333-0001

1 INST FOR ADVNCD TCHNLGY
THE UNIV OF TEXAS AT AUSTIN
PO BOX 202797
AUSTIN TX 78720-2797

1 DARPA
B KASPAR
3701 N FAIRFAX DR
ARLINGTON VA 22203-1714

1 NAVAL SURFACE WARFARE CTR
CODE B07 J PENNELLA
17320 DAHLGREN RD
BLDG 1470 RM 1101
DAHLGREN VA 22448-5100

1 US MILITARY ACADEMY
MATH SCI CTR OF EXCELLENCE
DEPT OF MATHEMATICAL SCI
MAJ M D PHILLIPS
THAYER HALL
WEST POINT NY 10996-1786

NO. OF
COPIES ORGANIZATION

1 DIRECTOR
US ARMY RESEARCH LAB
AMSRL D
R W WHALIN
2800 POWDER MILL RD
ADELPHI MD 20783-1145

1 DIRECTOR
US ARMY RESEARCH LAB
AMSRL DD
J J ROCCHIO
2800 POWDER MILL RD
ADELPHI MD 20783-1145

1 DIRECTOR
US ARMY RESEARCH LAB
AMSRL CS AS (RECORDS MGMT)
2800 POWDER MILL RD
ADELPHI MD 20783-1145

3 DIRECTOR
US ARMY RESEARCH LAB
AMSRL CI LL
2800 POWDER MILL RD
ADELPHI MD 20783-1145

ABERDEEN PROVING GROUND

4 DIR USARL
AMSRL CI LP (305)

NO. OF
COPIES ORGANIZATION

ABERDEEN PROVING GROUND

20 DIR USARL
 AMSRL IS T
 J BRAND

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
<small>Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.</small>				
1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE January 1999		3. REPORT TYPE AND DATES COVERED Final, January - November 1997
4. TITLE AND SUBTITLE A Proposed Modeling Protocol for Evaluating Information Attacks			5. FUNDING NUMBERS 2182040 6U-6U01 P611102.H48 *TEP10 BFTA0 S18129	
6. AUTHOR(S) John Brand II				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) U.S. Army Research Laboratory ATTN: AMSRL-IS-T Aberdeen Proving Ground, MD 21005-5067			8. PERFORMING ORGANIZATION REPORT NUMBER ARL-TN-112	
9. SPONSORING/MONITORING AGENCY NAMES(S) AND ADDRESS(ES)			10. SPONSORING/MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES				
12a. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.			12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) The essence of an information attack is to alter, either by intrusion into and manipulation of a database or by deception, the scenario under which a target mind or organization evaluates and selects future courses of action. The aim is to influence the actions of the target. The method is alteration of the perceived desirability or expected payoff of specific courses of action. This alteration of the information in possession of the target can be described as alteration of the perceived reality under which the target operates. Probable success by an attacker in altering the target's perceived behavior, given a successful manipulation of the target's information, has, in the past, been subjective. A modeling protocol based on the use of game theory is proposed that may, in certain cases, allow optimization of the scenario, or reality, imposed on the target to force the choice of a desired course of action. It should also allow a quantitative estimate of the likelihood of the target's adopting a given course of action. This tool can be used to estimate friendly susceptibility to information attack.				
14. SUBJECT TERMS information warfare, game theory, linear programming			15. NUMBER OF PAGES 30	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT SAR	

INTENTIONALLY LEFT BLANK.

USER EVALUATION SHEET/CHANGE OF ADDRESS

This Laboratory undertakes a continuing effort to improve the quality of the reports it publishes. Your comments/answers to the items/questions below will aid us in our efforts.

1. ARL Report Number/Author ARL-TN-112 (Brand) Date of Report January 1999

2. Date Report Received _____

3. Does this report satisfy a need? (Comment on purpose, related project, or other area of interest for which the report will be used.) _____

4. Specifically, how is the report being used? (Information source, design data, procedure, source of ideas, etc.) _____

5. Has the information in this report led to any quantitative savings as far as man-hours or dollars saved, operating costs avoided, or efficiencies achieved, etc? If so, please elaborate. _____

6. General Comments. What do you think should be changed to improve future reports? (Indicate changes to organization, technical content, format, etc.) _____

CURRENT
ADDRESS

Organization

Name

E-mail Name

Street or P.O. Box No.

City, State, Zip Code

7. If indicating a Change of Address or Address Correction, please provide the Current or Correct address above and the Old or Incorrect address below.

OLD
ADDRESS

Organization

Name

Street or P.O. Box No.

City, State, Zip Code

(Remove this sheet, fold as indicated, tape closed, and mail.)
(DO NOT STAPLE)